Defending Non-Bayesian Learning against Adversarial Attacks *

Lili Su Nitin H. Vaidya
Department of Electrical and Computer Engineering
University of Illinois at Urbana-Champaign
{lilisu3, nhv}@illinois.edu

Contact author and e-mail: Lili Su (lilisu3@illinois.edu)

arXiv:1606.08883v1 [cs.DC] 28 Jun 2016

Abstract

This paper addresses the problem of non-Bayesian learning over multi-agent networks, where agents repeatedly collect partially informative observations about an *unknown* state of the world, and try to collaboratively learn the true state. We focus on the impact of the adversarial agents on the performance of consensus-based non-Bayesian learning, where non-faulty agents combine local learning updates with consensus primitives. In particular, we consider the scenario where an unknown subset of agents suffer Byzantine faults – agents suffering Byzantine faults behave arbitrarily.

We propose two learning rules.

- In our first update rule, each agent updates its local beliefs as (up to normalization) the product of (1) the likelihood of the *cumulative* private signals and (2) the weighted geometric average of the beliefs of its incoming neighbors and itself. Under reasonable assumptions on the underlying network structure and the global identifiability of the network, we show that all the non-faulty agents asymptotically agree on the true state almost surely. For the case when every agent is failure-free, we show that (with high probability) each agent's beliefs on the wrong hypotheses decrease at rate $O(\exp(-Ct^2))$, where t is the number of iterations, and C is a constant.
- In general when agents may be adversarial, network identifiability condition specified for the above learning rule scales poorly in the number of state candidates m. In addition, the computation complexity per agent per iteration of this learning rule is forbiddingly high. Thus, we propose a modification of our first learning rule, whose complexity per iteration per agent is $O(m^2 n \log n)$, where n is the number of agents in the network. We show that this modified learning rule works under a much weaker network identifiability condition. In addition, this new condition is independent of m.

* This research is supported in part by National Science Foundation award NSF 1421918. Any opinions, findings, and conclusions or recommendations expressed here are those of the authors and do not necessarily reflect the views of the funding agencies or the U.S. government.

1 Introduction

Decentralized hypothesis testing (learning) has received significant amount of attention [1,2,3,5,6,7,8]. The traditional decentralized detection framework consists of a collection of spatially distributed sensors and a fusion center [5,6,7]. The sensors independently collect *noisy* observations of the environment state, and send only *summary* of the private observations to the fusion center, where a final decision is made. In the case when the sensors directly send all the private observations, the detection problem can be solved using a centralized scheme. The above framework does not scale well, since each sensor needs to be connected to the fusion center and full reliability of the fusion center is required, which may not be practical as the system scales.

Distributed hypothesis testing in the *absence* of fusion center is considered in [2,22,23,21]. In particular, Gale and Kariv [2] studied the distributed hypothesis testing problem in the context of social learning, where fully Bayesian belief update rule is studied. Bayesian update rule is impractical in many applications due to memory and computation constraints of each agent.

To avoid the complexity of Bayesian learning, a non-Bayesian learning framework that combines local Bayesian learning with distributed consensus was proposed by Jadbabaie et al. [3], and has attracted much attention [10,14,15,16,11,18,17,13]. Jadbabaie et al. [3] considered the general setting where external signals are observed during each iteration of the algorithm execution. Specifically, the belief of each agent is repeatedly updated as the arithmetic mean of its local Bayesian update and the beliefs of its neighbors – combining iterative consensus algorithm with local Bayesian update. It is shown [3] that, under this learning rule, each agent learns the true state almost surely. The publication of [3] has inspired significant efforts in designing and analyzing non-Bayesian learning rules with a particular focus on refining the fusion strategies and analyzing the (asymptotic and/or finite time) convergence rates of the refined algorithms [10,14,15,16,11,18,17,13]. In this paper we are particularly interested in the log-linear form of the update rule, in which, essentially, each agent updates its belief as the geometric average of the local Bayesian update and its neighbors' beliefs [15,10,14,16,11,18,17,13]. The log-linear form (geometric averaging) update rule is shown to converge exponentially fast [10,16]. Taking an axiomatic approach, the geometric averaging fusion is proved to be optimal [13]. An optimization-based interpretation of this rule is presented in [16], using dual averaging method with properly chosen proximal functions. Finite-time convergence rates are investigated independently in [14,11,17]. Both [14] and [18] consider time-varying networks, with slightly different network models. Specifically, [14] assumes that the union of every consecutive B networks is strongly connected, while [18] considers random networks. In this paper, we consider static networks for ease of exposition, although we believe that our results can be easily generalized to time-varying networks.

The prior work implicitly assumes that the networked agents are reliable in the sense that they correctly follow the specified learning rules. However, in some practical multi-agent networks, this assumption may not hold. For example, in social networks, it is possible that some agents are adversarial, and try to prevent the true state from being learned by the good agents. Thus, this paper focuses on the fault-tolerant version the non-Bayesian framework proposed in [3]. In particular, we assume that an unknown subset of agents may suffer Byzantine faults.

An agent suffering Byzantine fault may not follow the pre-specified algorithms/protocols, and *misbehave arbitrarily*. For instance, a faulty agent may lie to other agents (possibly non-consistently) about its own estimates. In addition, a faulty agent is assumed to have a complete knowledge of the system, including the network topology, the local functions of all the non-faulty agents, the algorithm specification of the non-faulty agents, the execution of the algorithm, the local estimates of all the non-faulty agents, and contents of messages the other agents send to each other. Also,

the faulty agents can potentially collaborate with each other to prevent the non-faulty agents from achieving their goal. An alternative fault model, where some agents may unexpectedly cease computing and communicate with each other asynchronously, is considered in our companion work [34]. The Byzantine fault-tolerance problem was introduced by Pease et al. [24] and has attracted intensive attention from researchers [25,26,27,30,28,31]. Our goal is to design algorithms that enable all the non-faulty agents to learn the underlying true state.

The existing non-Bayesian learning algorithms [10,11,13,14,15,16,17,18] are not robust to Byzantine agents, since the malicious messages sent by the Byzantine agents are indiscriminatingly utilized in the local belief updates. On the other hand, the incorporation of Byzantine consensus is non-trivial, since (i) the *effective* communication networks are *dependent* on the all the random local observations, making it non-trivial to adapt analysis of previous algorithms to our setting; and (ii) the problem of identifying tight topological condition for reaching Byzantine multi-dimensional consensus iteratively is open, making it challenging to identify the minimal detectability condition on the networked agents to learn the true environmental state.

Contributions: Our contributions are two-fold.

- We first propose an update rule wherein each agent iteratively updates its local beliefs as (up to normalization) the product of (1) the likelihood of the *cumulative* private signals and (2) the weighted geometric average of the beliefs of its incoming neighbors and itself (using iterative Byzantine multi-dimensional consensus). In contrast to the existing algorithms [14,11], where only the *current* private signal is used in the update, our proposed algorithm relies on the *cumulative* private signals. Under reasonable assumptions on the underlying network structure and the global identifiability of the network, we show that all the non-faulty agents asymptotically agree on the true state almost surely. In addition, for the special case when every agent is guaranteed to be failure-free, we show that (with high probability) each agent's beliefs on the wrong hypotheses decrease at rate $O(\exp(-Ct^2))$, where t is the number of iterations, and C is a constant. Thus, our proposed rule may be of independent interest for the failure-free setting considered in [10,11,13,14,15,16,17,18].
- The local computation complexity per agent of the first learning rule is high due to the adoption of multi-dimensional consensus primitives. More importantly, the network identifiability condition used for that learning rule scales poorly in the number of possible states m. Thus, we propose a modification of our first learning rule, whose complexity per iteration per agent is $O(m^2 n \log n)$, where n is the number of agents in the network. We show that this modified learning rule works under a much weaker global identifiability condition, which is independent of m. We cast the general m–ary hypothesis testing problem into a collection of binary hypothesis testing sub-problems.

Outline: The rest of the paper is organized as follows. Section 2 presents the problem formulation. Section 3 briefly reviews existing results on vector Byzantine consensus, and matrix representation of the state evolution. Our first algorithm and its correctness analysis are presented in Section 4. Section 5 demonstrates the above learning rule in the special case when $f = 0$, and presents a finite-time analysis.The modified learning rule and its correctness analysis are summarized in Section 6. Section 7 concludes the paper, and discusses possible extensions.

2 Problem Formulation

Network Model: Our network model is similar to the model used in [4,30]. We consider a synchronous system. A collection of n agents (also referred as *nodes*) are connected by a *directed* network $G(\mathcal{V}, \mathcal{E})$,

where $\mathcal{V} = \{1, \ldots, n\}$ and $\mathcal{E}$ is the collection of *directed* edges. For each $i \in \mathcal{V}$, let $\mathcal{I}_i$ denote the set of incoming neighbors of agent i. In any execution, up to f agents suffer Byzantine faults. For a given execution, let $\mathcal{F}$ denote the set of Byzantine agents, and $\mathcal{N}$ denote the set of non-faulty agents. Throughout this paper, we assume that f satisfies the condition implicitly imposed by the given topology conditions mentioned later. We assume that each non-faulty agent knows f, but does not know the *actual* number of faulty agents $|\mathcal{F}|$. [1] Possible misbehavior of faulty agents includes sending incorrect and mismatching (or inconsistent) messages. The Byzantine agents are also assumed to have complete knowledge of system, including the network topology, underlying running algorithm, the states or even the entire history. The faulty agents may collaborate with each other adaptively [12]. Note that $|\mathcal{F}| \leq f$ and $|\mathcal{N}| \geq n - f$ since at most f agents may fail. (As noted earlier, although we assume a static network topology, our results can be easily generalized to time-varying networks.)

Throughout this paper, we use the terms *agent* and *node* interchangeably.

Observation Model: Our observation model is identical the model used in [3,11,18]. Let $\Theta = \{\theta_1, \theta_2, \ldots, \theta_m\}$ denote a set of m environmental states, which we call *hypotheses*. In the t-th iteration, each agent *independently* obtains a private signal about the environmental state θ^*, which is initially unknown to every agent in the network. Each agent i knows the structure of its private signal, which is represented by a collection of parameterized marginal distributions $\mathcal{D}^i = \{\ell_i(w_i|\theta) | \, \theta \in \Theta, \, w_i \in \mathcal{S}_i\}$, where $\ell_i(\cdot|\theta)$ is the distribution of private signal when θ is the true state, and $\mathcal{S}_i$ is the finite private signal space. For each $\theta \in \Theta$, and each $i \in \mathcal{V}$, the support of $\ell_i(\cdot|\theta)$ is the whole signal space, i.e., $\ell_i(w_i|\theta) > 0$, $\forall \, w_i \in \mathcal{S}_i$ and $\forall \, \theta \in \Theta$. Let s_t^i be the private signal observed by agent i in iteration t, and let $\mathbf{s}_t = \{s_t^1, s_t^2, \ldots, s_t^n\}$ be the signal profile at time t (i.e., signals observed by the agents in iteration t). Given an environmental state θ, the signal profile $\mathbf{s}_t$ is generated according to the joint distribution $\ell_1(s_t^1|\theta) \times \ell_2(s_t^2|\theta) \times \cdots \times \ell_n(s_t^n|\theta)$. In addition, let $s_{1,t}^i$ be the signal history up to time t for agent $i = 1, \cdots, n$, and let $\mathbf{s}_{1,t} = \{s_{1,t}^1, s_{1,t}^2, \ldots, s_{1,t}^n\}$ be the signal profile history up to time t.

3 Byzantine Consensus

In this section, we briefly review relevant exsting results on Byzantine consensus. Byzantine consensus has attracted significant amount of attention [25,26,29,27,30,28,31]. While the past work mostly focus on scalar inputs, the more general vector (or multi-dimensional) inputs have been studied recently [31,29,28]. Complete communication networks are considered in [31,29], where tight conditions on the number of agents are identified. Incomplete communication networks are studied in [28]. Closer to the non-Bayesian learning problem is the class of *iterative approximate Byzantine consensus algorithms*, where each agent is only allowed to exchange information about its state with its neighbors. In particular, our learning algorithms build upon *Byz-Iter* algorithm proposed in [28] and a simple algorithm proposed in [30] for iterative Byzantine consensus with vector inputs and scalar inputs, respectively, in incomplete networks. A matrix representation of the non-faulty agents' states evolution under *Byz-Iter* algorithm is provided by [28], which also captures the dynamics of the simple algorithm with scalar inputs in [30]. To make this paper self-contained, in this section, we briefly review the algorithm *Byz-Iter* and its matrix representation.

3.1 Algorithm *Byz-Iter* [28]

Algorithm *Byz-Iter* is based on Tverberg's Theorem [32].

[1] This is because the upper bound f can be learned via long-time performance statistics, whereas, the actual size of $\mathcal{F}$ varies across executions, and may be impossible to be predicted in some applications.

Theorem 1. *[32] Let f be a nonnegative integer. Let Y be a multiset containing vectors from $\mathbb{R}^m$ such that $|Y| \geq (m+1)f+1$. There exists a partition $Y_1, Y_2, \cdots, Y_{f+1}$ of Y such that Y_i is nonempty for $1 \leq i \leq f+1$, and the intersection of the convex hulls of Y_i's are nonempty, i.e., $\cap_{i=1}^{f+1} \mathsf{Conv}(Y_i) \neq \emptyset$, where $\mathsf{Conv}(Y_i)$ is the convex hull of Y_i for $i = 1, \cdots, f+1$.*

The proper partition in Theorem 1, and the points in $\cap_{i=1}^{f+1} \mathsf{Conv}(Y_i)$, are referred as *Tverberg partition of Y* and *Tverberg points of Y*, respectively.

For convenience of presenting our algorithm in Section 4, we present *Byz-Iter* (described in Algorithm 2) below using *One-Iter* (described in Algorithm 1) as a primitive. The parameter $\mathbf{x}^i$ passed to *One-Iter* at agent i, and $\mathbf{y}^i$ returned by *One-Iter* are both m-dimensional vectors. Let $\mathbf{v}^i$ be the state of agent i that will be iteratively updated, with $\mathbf{v}_t^i$ being the state at the end of iteration t and $\mathbf{v}_0^i$ being the input of agent i. In each iteration $t \geq 1$, a non-faulty agent performs the steps in *One-Iter*. In particular, in the message receiving step, if a message is not received from some neighbor, that neighbor must be faulty, as the system is synchronous. In this case, the missing message values are set to some default value. Faulty agents may deviate from the algorithm specification arbitrarily. In *Byz-Iter*, the value returned by *One-Iter* at agent i is assigned to $\mathbf{v}_t^i$.

Algorithm 1: Algorithm *One-Iter* with input $\mathbf{x}^i$ at agent i

$Z^i \leftarrow \emptyset$;
Transmit $\mathbf{x}^i$ on all outgoing links;
Receive messages on all incoming links. % These message values form a multiset R^i of size $|\mathcal{I}_i|$.%
for *every* $C \subseteq R^i \cup \{\mathbf{x}^i\}$ *such that* $|C| = (m+1)f+1$ **do**
  add to Z^i a *Tverberg point* of multiset C
end
Compute $\mathbf{y}^i$ as follows: $\mathbf{y}^i \leftarrow \frac{1}{1+|Z^i|}\left(\mathbf{x}^i + \sum_{\mathbf{z} \in Z^i} \mathbf{z}\right)$;
Return $\mathbf{y}^i$;

Algorithm 2: Algorithm *Byz-Iter* [28]: t-th iteration at agent i

$\mathbf{v}_t^i \leftarrow$ *One-Iter*$(\mathbf{v}_{t-1}^i)$;

Remark 1. Note that for each agent $i \in \mathcal{N}$, the computation complexity per iteration is

$$\Omega\left(\binom{|R^i \cup \{\mathbf{x}^i\}|}{(m+1)f+1}\right) = \Omega\left(\binom{|\mathcal{I}_i|+1}{(m+1)f+1}\right).$$

In the worst case, $||\mathcal{I}_i|+1| = n$, and

$$\Omega\left(\binom{|\mathcal{I}_i|+1}{(m+1)f+1}\right) = \Omega\left(\binom{n}{(m+1)f+1}\right) = \Omega\left(\left(\frac{n}{\mathrm{e}}\right)^{(m+1)f+1}\right).$$

Since our first learning rule is based on Algorithm *Byz-Iter*, the computation complexity of our first proposed algorithm is also high. Nevertheless, our first learning rule contains our main algorithmic

ideas. More importantly, this learning rule can be modified such that the computation complexity per iteration per agent is $O(m^2 n \log n)$. Specifically, the modified learning rule adopts the scalar Byzantine consensus instead of the m–dimensional consensus. This modified learning rule is optimal in the sense that it works under minimal network identifiability requirements.

3.2 Correctness of Algorithm *Byz-Iter*

We briefly summarize the aspects of correctness proof of Algorithm 2 from [28] that are necessary for our subsequent discussion. By using the Tverberg points in the update of $\mathbf{v}_t^i$ above, effectively, the extreme message values (that may potentially be sent by faulty agents) are trimmed away. Informally speaking, trimming certain messages can be viewed as ignoring (or removing) incoming links that carry the outliers. [28] shows that the effective communication network thus obtained can be characterized by a "reduced graph" of $G(\mathcal{V}, \mathcal{E})$, defined below. It is important to note that the non-faulty agents **do not** know the identity of the faulty agents.

Definition 1 (m–dimensional reduced graph). *An m–dimensional reduced graph $\mathcal{H}(\mathcal{N}, \mathcal{E}_{\mathcal{F}})$ of $G(\mathcal{V}, \mathcal{E})$ is obtained by (i) removing all faulty nodes $\mathcal{F}$, and all the links incident on the faulty nodes $\mathcal{F}$; and (ii) for each non-faulty node (nodes in $\mathcal{N}$), removing up to mf additional incoming links.*

Definition 2. *A source component in any given m–dimensional reduced graph is a strongly connected component (of that reduced graph), which does not have any incoming links from outside that component.*

It turns out that the effective communication network is potentially time-varying (partly) due to time-varying behavior of faulty nodes. Assumption 1 below states a condition that is sufficient for reaching approximate Byzantine vector consensus using Algorithm 1 [28].

Assumption 1 *Every m–dimensional reduced graph of $G(\mathcal{V}, \mathcal{E})$ contains a unique source component.*

Let $\mathcal{C}_m$ be the set of all the m–dimensional reduced graph of $G(\mathcal{V}, \mathcal{E})$. Define $\chi_m \triangleq |\mathcal{C}_m|$. Since $G(\mathcal{V}, \mathcal{E})$ is finite, we have $\chi_m < \infty$. Let $\mathcal{H}_m \in \mathcal{C}_m$ be an m–dimensional reduced graph of $G(\mathcal{V}, \mathcal{E})$ with source component $\mathcal{S}_{\mathcal{H}_m}$. Define

$$\gamma_m \triangleq \min_{\mathcal{H}_m \in \mathcal{C}_m} |\mathcal{S}_{\mathcal{H}_m}|, \tag{1}$$

i.e., γ_m is the minimum source component size among all the m–dimensional reduced graphs. Note that $\gamma_m \geq 1$ if Assumption 1 holds for a given m.

Theorem 2. *[28] Suppose Assumption 1 holds for a given $m \geq 1$. Under Algorithm* Byz-Iter*, all the non-faulty agents (agents in $\mathcal{N}$) reach consensus asymptotically, i.e., $\lim_{t\to\infty} |\mathbf{v}_t^i - \mathbf{v}_t^j| = 0, \forall\, i, j \in \mathcal{N}$.*

The proof of Theorem 2 relies crucially on a matrix representation of the state evolution.

3.3 Matrix Representation [28]

Let $|\mathcal{F}| = \phi$ (thus, $0 \leq \phi \leq f$). Without loss of generality, assume that agents 1 through $n - \phi$ are non-faulty, and agents $n - \phi + 1$ to n are Byzantine.

Lemma 1. *[28] Suppose Assumption 1 holds for a given $m \geq 1$. The state updates performed by the non-faulty agents in the t–th iteration ($t \geq 1$) can be expressed as*

$$\mathbf{v}_t^i = \sum_{j=1}^{n-\phi} \mathbf{A}_{ij}[t]\mathbf{v}_{t-1}^j, \tag{2}$$

where $\mathbf{A}[t] \in \mathbb{R}^{(n-\phi)\times(n-\phi)}$ is a row stochastic *matrix for which there exists an m–dimensional reduced graph $\mathcal{H}_m[t]$ with adjacency matrix $\mathbf{H}_m[t]$ such that $\mathbf{A}[t] \geq \beta_m \mathbf{H}_m[t]$, where $0 < \beta_m \leq 1$ is a constant that depends only on $G(\mathcal{V}, \mathcal{E})$.*

Let $\mathbf{\Phi}(t, r) \triangleq \mathbf{A}[t] \cdots \mathbf{A}[r]$ for $1 \leq r \leq t+1$. By convention, $\mathbf{\Phi}(t, t) = \mathbf{A}[t]$ and $\mathbf{\Phi}(t, t+1) = \mathbf{I}$. Note that $\mathbf{\Phi}(t, r)$ is a backward product. Using prior work on coefficients of ergodicity [9], under Assumption 1, it has been shown [28,19] that

$$\lim_{t \geq r,\ t \to \infty} \mathbf{\Phi}(t, r) = \mathbf{1}\pi(r), \tag{3}$$

where $\pi(r) \in \mathbb{R}^{n-\phi}$ is a row stochastic vector, and $\mathbf{1}$ is the column vector with each entry being 1. Recall that χ_m is the total number of m–dimensional reduced graphs of $G(\mathcal{V}, \mathcal{E})$, and β_m is defined in Lemma 1, and $\phi \triangleq |\mathcal{F}|$. The convergence rate in (3) is exponential.

Theorem 3. *[28] For all $t \geq r \geq 1$, it holds that $|\mathbf{\Phi}_{ij}(t, r) - \pi_j(r)| \leq (1 - \beta_m^{\nu})^{\lceil \frac{t-r+1}{\nu} \rceil}$, where $\nu \triangleq \chi_m(n - \phi)$.*

Recall that γ_m is defined in (1). The next lemma is a consequence of the results in [28].

Lemma 2. *[28] For any $r \geq 1$, there exists a reduced graph $\mathcal{H}[r]$ with source component $\mathcal{S}_r$ such that $\pi_i(r) \geq \beta_m^{\chi_m(n-\phi)}$ for each $i \in \mathcal{S}_r$. In addition, $|\mathcal{S}_r| \geq \gamma_m$.*

3.4 Tight Topological Condition for Scalar Iterative Byzantine Consensus

The above analysis shows that Assumption 1 is sufficient for achieving Byzantine consensus iteratively. For the special case when $m = 1$,(i.e., the inputs provided at individual non-faulty agents are scalars) it has been shown [30] that Assumption 1 is also necessary.

Theorem 4. *[30] For scalar inputs, iterative approximate Byzantine consensus is achievable among non-faulty agents if and only if every 1-dimensional reduced graph of $G(\mathcal{V}, \mathcal{E})$ contains only one source component.*

Moreover, the following simple algorithm (Algorithm 3) works under Assumption 1 when $m = 1$.

In addition, it has been show that the dynamic of the non-faulty agents states admits the same matrix representation as in Subsection 3.3 with the reduced graph being 1–dimensional reduced graph defined in Definition 1.

With the above background on Byzantine vector consensus, we are now ready to present our first algorithm and its analysis.

4 Byzantine Fault-Tolerant Non-Bayesian Learning (BFL)

In this section, we present our first learning rule, named Byzantine Fault-Tolerant Non-Bayesian Learning (BFL). In BFL, each agent i maintains a belief vector $\mu^i \in \mathbb{R}^m$, which is a distribution over the set Θ, with $\mu^i(\theta)$ being the probability with which the agent i *believes* that θ is the true

Algorithm 3: Algorithm Scalar Byzantine Consensus: iteration $t \geq 1$ [30]

1 Transmit $v^i[t-1]$ on all outgoing links;
2 Receive messages on all incoming links. % These message values $w_j[t]$ for each $j \in \mathcal{I}_i$ form a multiset $R^i[t]$ of size $|\mathcal{I}_i|$. %
3 Sort the received values $w_j[t]$ for each $j \in \mathcal{I}_i$ in a non-decreasing order;
4 Remove the largest f values and the smallest f values. % Denote the set of indices of incoming neighbors whose values have not been removed at iteration t by $\mathcal{I}_i^*[t]$.%
5 Update v^i as follows: $v^i[t] \leftarrow \frac{\sum_{j \in \mathcal{I}_i^*[t]} w_j[t] + v^i[t-1]}{1+|\mathcal{I}_i^*[t]|}$;

environmental state. Since no signals are observed before the execution of an algorithm, the belief μ^i is often initially set to be uniform over the set Θ, i.e., $\big(\mu_0^i(\theta_1), \mu_0^i(\theta_1), \ldots, \mu_0^i(\theta_m)\big)^T = \big(\frac{1}{m}, \ldots, \frac{1}{m}\big)^T$. Recall that θ^* is the true environmental state. We say the networked agents collaboratively learn θ^* if for every non-faulty agent $i \in \mathcal{N}$,

$$\lim_{t \to \infty} \mu_t^i(\theta^*) \;=\; 1 \;\; a.s. \tag{4}$$

where $a.s.$ denotes *almost surely*.

BFL is a modified version of the geometric averaging update rule that has been investigated in previous work [14,15,11,17]. In particular, we modify the averaging rule to take into account Byzantine faults. More importantly, in each iteration, we use the likelihood of the *cumulative* local observations (instead of the likelihood of the *current* observation only) to update the local beliefs.

For $t \geq 1$, the steps to be performed by agent i in the t–th iteration are listed below. Note that faulty agents can deviate from the algorithm specification. The algorithm below uses *One-Iter* presented in the previous section as a primitive. Recall that $s_{1,t}^i$ is the cumulative local observations up to iteration t. Since the observations are *i.i.d.*, it holds that $\ell_i(s_{1,t}^i|\theta) = \prod_{r=1}^t \ell_i(s_r^i|\theta)$. So $\ell_i(s_{1,t}^i|\theta)$ can be computed iteratively in Algorithm 4.

Algorithm 4: BFL: Iteration $t \geq 1$ at agent i

1 $\eta_t^i \leftarrow$ *One-Iter*$(\log \mu_{t-1}^i)$;
2 Observe s_t^i;
3 **for** $\theta \in \Theta$ **do**
4 $\quad \ell_i(s_{1,t}^i|\theta) \leftarrow \ell_i(s_t^i|\theta)\, \ell_i(s_{1,t-1}^i|\theta)$;
5 $\quad \mu_t^i(\theta) \leftarrow \frac{\ell_i(s_{1,t}^i|\theta) \exp\big(\eta_t^i(\theta)\big)}{\sum_{p=1}^m \ell_i(s_{1,t}^i|\theta_p) \exp\big(\eta_t^i(\theta_p)\big)}$;
6 **end**

The main difference of Algorithm 4 with respect to the algorithms in [14,15,11,17] is that (i) our algorithm uses a Byzantine consensus iteration as a primitive (in line 1), and (ii) $\ell_i(s_{1,t}^i|\theta)$ used in line 5 is the likelihood for observations from iteration 1 to t (the previous algorithms instead use $\ell_i(s_t^i|\theta)$ here). Observe that the consensus step is being performed on log of the beliefs, with the result being stored as η_t^i (in line 1) and used in line 4 to compute the new beliefs.

Recalling the matrix representation of the *Byz-Iter* algorithm as per Lemma 1, we can write the following equivalent representation of line 1 of Algorithm 4.

$$\eta_t^i(\theta) = \sum_{j=1}^{n-\phi} \mathbf{A}_{ij}[t] \log \mu_{t-1}^j(\theta) = \log \prod_{j=1}^{n-\phi} \mu_{t-1}^j(\theta)^{\mathbf{A}_{ij}[t]}, \qquad \forall \theta \in \Theta. \tag{5}$$

where $\mathbf{A}[t]$ is a row stochastic matrix whose properties are specified in Lemma 1. Note that $\mu_t^i(\theta)$ is **random** for each $i \in \mathcal{N}$ and $t \geq 1$, as it is updated according to local random observations. Since the consensus is performed over $\log \mu_t^i \in \mathbb{R}^m$, the update matrix $\mathbf{A}[t]$ is also **random**. In particular, for each $t \geq 1$, matrix $\mathbf{A}[t]$ is dependent on *all the cumulative observations over the network* up to iteration t. This dependency makes it non-trivial to adapt analysis from previous algorithms to our setting. In addition, adopting the local cumulative observation likelihood makes the analysis with Byzantine faults easier.

4.1 Identifiability

In the absence of agent failures [3], for the networked agents to detect the true hypothesis θ^*, it is sufficient to assume that $G(\mathcal{V}, \mathcal{E})$ is strongly connected, and that θ^* is globally identifiable. That is, for any $\theta \neq \theta^*$, there exists a node $j \in \mathcal{V}$ such that the Kullback-Leiber divergence between the true marginal $\ell_j(\cdot|\theta^*)$ and the marginal $\ell_j(\cdot|\theta)$, denoted by $D\left(\ell_j(\cdot|\theta^*)||\ell_j(\cdot|\theta)\right)$, is nonzero; equivalently,

$$\sum_{j \in \mathcal{V}} D\left(\ell_j(\cdot|\theta^*)||\ell_j(\cdot|\theta)\right) \;\neq\; 0, \tag{6}$$

where $D\left(\ell_j(\cdot|\theta^*)||\ell_j(\cdot|\theta)\right)$ is defined as

$$D\left(\ell_j(\cdot|\theta^*)||\ell_j(\cdot|\theta)\right) \triangleq \sum_{w_j \in \mathcal{S}_j} \ell_j(w_j|\theta^*) \log \frac{\ell_j(w_j|\theta^*)}{\ell_j(w_j|\theta)}. \tag{7}$$

Since θ^* may change from execution to execution, (6) is required to hold for any choice of θ^*. Intuitively speaking, if any pair of states θ_1 and θ_2 can be distinguished by at least one agent in the network, then sufficient exchange of local beliefs over strongly connected network will enable every agent distinguish θ_1 and θ_2. However, in the presence of Byzantine agents, a stronger global identifiability condition is required. The following assumption builds upon Assumption 1.

Assumption 2 *Suppose that Assumption 1 holds for $m = |\Theta|$. For any $\theta \neq \theta^*$, and for any m–dimensional reduced graph $\mathcal{H}$ of $G(\mathcal{V}, \mathcal{E})$ with $\mathcal{S}_\mathcal{H}$ denoting the unique source component, the following holds*

$$\sum_{j \in \mathcal{S}_\mathcal{H}} D\left(\ell_j(\cdot|\theta^*) \parallel \ell_j(\cdot|\theta)\right) \;\neq\; 0. \tag{8}$$

In contrast to (6), where the summation is taken over all the agents in the network, in (8), the summation is taken over agents in the source component only. Intuitively, the condition imposed by Assumption 2 is that all the agents in the source component can detect the true state θ^* collaboratively. If iterative consensus is achieved, the accurate belief can be propagated from the source component to every other non-faulty agent in the network.

Remark 2. We will show later that when Assumption 2 holds, BFL algorithm enables all the non-faulty agents concentrate their beliefs on the true state θ^* almost surely. That is, Assumption 2 is a sufficient condition for a consensus-based non-Bayesian learning algorithm to exist. However, Assumption 2 is not necessary, observing that Assumption 1 (upon which Assumption 2 builds) is not necessary for m-dimensional Byzantine consensus algorithms to exist. As illustrated by our second learning rule (described later), the adoption of m-dimensional Byzantine consensus primitives is *not necessary*. Nevertheless, BFL contains our main algorithmic and analytical ideas. In addition, BFL provides an alternative learning rule for the failure-free setting (where no fault-tolerant consensus primitives are needed).

4.2 Convergence Results

Our proof parallels the structure of a proof in [14], but with some key differences to take into account our update rule for the belief vector.

For any $\theta_1, \theta_2 \in \Theta$, and any $i \in \mathcal{V}$, define $\boldsymbol{\psi}_t^i(\theta_1, \theta_2)$ and $\mathcal{L}_t(\theta_1, \theta_2)$ as follows

$$\boldsymbol{\psi}_t^i(\theta_1, \theta_2) \triangleq \log \frac{\mu_t^i(\theta_1)}{\mu_t^i(\theta_2)}, \quad \mathcal{L}_t^i(\theta_1, \theta_2) \triangleq \log \frac{\ell_i(s_t^i|\theta_1)}{\ell_i(s_t^i|\theta_2)}. \tag{9}$$

To show Algorithm 4 solves (4), we will show that $\boldsymbol{\psi}_t^i(\theta, \theta^*) \xrightarrow{\text{a.s.}} -\infty$ for $\theta \neq \theta^*$, which implies that $\mu_t^i(\theta) \xrightarrow{\text{a.s.}} 0$ for all $\theta \neq \theta^*$ and for all $i \in \mathcal{N}$, i.e., all non-faulty agents asymptotically concentrate their beliefs on the true hypothesis θ^*. We do this by investigating the dynamics of beliefs which is represented compactly in a matrix form.

For each $\theta \neq \theta^*$, and each $i \in \mathcal{N} = \{1, 2, \cdots, n-\phi\}$, we have

$$\begin{aligned}
\boldsymbol{\psi}_t^i(\theta, \theta^*) &= \log \frac{\mu_t^i(\theta)}{\mu_t^i(\theta^*)} \overset{(a)}{=} \log \left(\prod_{j=1}^{n-\phi} \left(\frac{\mu_{t-1}^j(\theta)}{\mu_{t-1}^j(\theta^*)} \right)^{\mathbf{A}_{ij}[t]} \times \frac{\ell_i(s_{1,t}^i|\theta)}{\ell_i(s_{1,t}^i|\theta^*)} \right) \\
&= \sum_{j=1}^{n-\phi} \mathbf{A}_{ij}[t] \log \frac{\mu_{t-1}^j(\theta)}{\mu_{t-1}^j(\theta^*)} + \log \frac{\ell_i(s_{1,t}^i|\theta)}{\ell_i(s_{1,t}^i|\theta^*)} \\
&= \sum_{j=1}^{n-\phi} \mathbf{A}_{ij}[t] \boldsymbol{\psi}_{t-1}^j(\theta, \theta^*) + \sum_{r=1}^{t} \mathcal{L}_r^i(\theta, \theta^*),
\end{aligned} \tag{10}$$

where equality (a) follows from (5) and the update of μ^i in Algorithm 4, and the last equality follows from (9) and the fact that the local observations are *i.i.d.* for each agent.

Let $\boldsymbol{\psi}_t(\theta, \theta^*) \in \mathbb{R}^{n-\phi}$ be the vector that stacks $\boldsymbol{\psi}_t^i(\theta, \theta^*)$, with the i–th entry being $\boldsymbol{\psi}_t^i(\theta, \theta^*)$ for all $i \in \mathcal{N}$. The evolution of $\boldsymbol{\psi}(\theta, \theta^*)$ can be compactly written as

$$\boldsymbol{\psi}_t(\theta, \theta^*) = \mathbf{A}[t] \boldsymbol{\psi}_{t-1}(\theta, \theta^*) + \sum_{r=1}^{t} \mathcal{L}_r(\theta, \theta^*). \tag{11}$$

Expanding (11), we get

$$\boldsymbol{\psi}_t(\theta, \theta^*) = \boldsymbol{\Phi}(t, 1) \boldsymbol{\psi}_0(\theta, \theta^*) + \sum_{r=1}^{t} \boldsymbol{\Phi}(t, r+1) \sum_{k=1}^{r} \mathcal{L}_k(\theta, \theta^*). \tag{12}$$

For each $\theta \in \Theta$ and $i \in \mathcal{V}$, define $H_i(\theta, \theta^*) \in \mathbb{R}^{n-\phi}$ as

$$\begin{aligned} H_i(\theta, \theta^*) &\triangleq \sum_{w_i \in \mathcal{S}_i} \ell_i(w_i|\theta^*) \log \frac{\ell_i(w_i \mid \theta)}{\ell_i(w_i \mid \theta^*)} \\ &= -D(\ell_i(\cdot|\theta^*) \parallel \ell_i(\cdot|\theta)) \quad \text{by (7)} \\ &\leq 0. \end{aligned} \tag{13}$$

Let $\mathcal{H} \in \mathcal{C}$ be an arbitrary reduced graph with source component $\mathcal{S}_{\mathcal{H}}$. Define C_0 and C_1 as

$$-C_0 \triangleq \min_{i \in \mathcal{V}} \min_{\theta_1, \theta_2 \in \Theta; \theta_1 \neq \theta_2} \min_{w_i \in \mathcal{S}_i} \left(\log \frac{\ell_i(w_i|\theta_1)}{\ell_i(w_i|\theta_2)} \right), \tag{14}$$

$$C_1 \triangleq \min_{\mathcal{H} \in \mathcal{C}} \min_{\theta, \theta^* \in \Theta; \theta \neq \theta^*} \sum_{i \in \mathcal{S}_{\mathcal{H}}} D(\ell_i(\cdot|\theta^*) \parallel \ell_i(\cdot|\theta)). \tag{15}$$

The constant C_0 serves as an universal upper bound on $|\log \frac{\ell_i(w_i|\theta_1)}{\ell_i(w_i|\theta_2)}|$ for all choices of θ_1 and θ_2, and for all signals. Intuitively, the constant C_1 is the minimal detection capability of the source component under Assumption 2.

Due to $|\Theta| = m < \infty$ and $|\mathcal{S}_i| < \infty$ for each $i \in \mathcal{N}$, we know that $C_0 < \infty$. Besides, it is easy to see that $-C_0 \leq 0$ (thus, $C_0 \geq 0$). In addition, under Assumption 2, we have $C_1 > 0$.

Now we present a key lemma for our main theorem.

Lemma 3. *Under Assumption 2, for any $\theta \neq \theta^*$, it holds that*

$$\frac{1}{t^2} \sum_{r=1}^{t} \left(\sum_{j=1}^{n-\phi} \mathbf{\Phi}_{ij}(t, r+1) \sum_{k=1}^{r} \mathcal{L}_k^j(\theta, \theta^*) - r \sum_{j=1}^{n-\phi} \pi_j(r+1) H_j(\theta, \theta^*) \right) \xrightarrow{\text{a.s.}} 0. \tag{16}$$

As it can be seen later, the proof of Lemma 3 is significantly different from the analogous lemma in [14].

Theorem 5. *When Assumption 2 holds, each non-faulty agent $i \in \mathcal{N}$ will concentrate its belief on the true hypothesis θ^* almost surely, i.e., $\mu_t^i(\theta) \xrightarrow{\text{a.s.}} 0$ for all $\theta \neq \theta^*$.*

Proof. Consider any $\theta \neq \theta^*$. Recall from (12) that

$$\begin{aligned} \boldsymbol{\psi}_t(\theta, \theta^*) &= \mathbf{\Phi}(t, 1) \boldsymbol{\psi}_0(\theta, \theta^*) + \sum_{r=1}^{t} \mathbf{\Phi}(t, r+1) \sum_{k=1}^{r} \mathcal{L}_k(\theta, \theta^*) \\ &= \sum_{r=1}^{t} \mathbf{\Phi}(t, r+1) \sum_{k=1}^{r} \mathcal{L}_k(\theta, \theta^*). \end{aligned}$$

The last equality holds as μ_0^i is uniform, and $\boldsymbol{\psi}_0^i(\theta, \theta^*) = 0$ for each $i \in \mathcal{N}$. Since the supports of $\ell_i(\cdot|\theta)$ and $\ell_i(\cdot|\theta^*)$ are the whole signal space $\mathcal{S}_i$ for each agent $i \in \mathcal{N}$, it holds that $\left| \frac{\ell_i(w_i|\theta)}{\ell_i(w_i|\theta^*)} \right| < \infty$ for each $w_i \in \mathcal{S}_i$, and

$$0 \geq H_i(\theta, \theta^*) \geq \min_{w_i \in \mathcal{S}_i} \left(\log \frac{\ell_i(w_i|\theta)}{\ell_i(w_i|\theta^*)} \right) \geq -C_0 > -\infty. \tag{17}$$

By (17), we know that $|\sum_{j=1}^{n-\phi}\pi_j(r+1)H_j(\theta,\theta^*)| \le C_0 < \infty$. Due to the finiteness of $\sum_{j=1}^{n-\phi}\pi_j(r+1)H_j(\theta,\theta^*)$, we are able to add and subtract $r\mathbf{1}\sum_{j=1}^{n-\phi}\pi_j(r+1)H_j(\theta,\theta^*)$ from (12). We get

$$\begin{aligned}\boldsymbol{\psi}_t(\theta,\theta^*) = &\sum_{r=1}^{t}\left(\boldsymbol{\Phi}(t,r+1)\sum_{k=1}^{r}\mathcal{L}_k(\theta,\theta^*) - r\mathbf{1}\sum_{j=1}^{n-\phi}\pi_j(r+1)H_j(\theta,\theta^*)\right)\\ &+\sum_{r=1}^{t}r\mathbf{1}\sum_{j=1}^{n-\phi}\pi_j(r+1)H_j(\theta,\theta^*).\end{aligned} \tag{18}$$

For each $i \in \mathcal{N}$, we have

$$\begin{aligned}\boldsymbol{\psi}_t^i(\theta,\theta^*) = &\sum_{r=1}^{t}\left(\sum_{j=1}^{n-\phi}\boldsymbol{\Phi}_{ij}(t,r+1)\sum_{k=1}^{r}\mathcal{L}_k^j(\theta,\theta^*) - r\sum_{j=1}^{n-\phi}\pi_j(r+1)H_j(\theta,\theta^*)\right)\\ &+\sum_{r=1}^{t}r\sum_{j=1}^{n-\phi}\pi_j(r+1)H_j(\theta,\theta^*).\end{aligned} \tag{19}$$

To show $\lim_{t\to\infty}\mu_t^i(\theta) \xrightarrow{\text{a.s.}} 0$ for $\theta \neq \theta^*$, it is enough to show $\boldsymbol{\psi}_t^i(\theta,\theta^*) \xrightarrow{\text{a.s.}} -\infty$. Our convergence proof has similar structure as the analysis in [14]. From Lemma 3, we know that

$$\frac{1}{t^2}\sum_{r=1}^{t}\left(\sum_{j=1}^{n-\phi}\boldsymbol{\Phi}_{ij}(t,r+1)\sum_{k=1}^{r}\mathcal{L}_k^j(\theta,\theta^*) - r\sum_{j=1}^{n-\phi}\pi_j(r+1)H_j(\theta,\theta^*)\right) \xrightarrow{\text{a.s.}} 0. \tag{20}$$

Next we show that the second term of the right hand side of (19) decreases quadratically in t.

$$\begin{aligned}\sum_{r=1}^{t}r\sum_{j=1}^{n-\phi}\pi_j(r+1)H_j(\theta,\theta^*) &\le \sum_{r=1}^{t}r\sum_{j\in\mathcal{S}_r}\pi_j(r+1)H_j(\theta,\theta^*) \qquad \text{by (13)}\\ &\le \sum_{r=1}^{t}r\beta^{\chi(n-\phi)}\sum_{j\in\mathcal{S}_r}H_j(\theta,\theta^*) \qquad \text{by Lemma 2}\\ &\le -\sum_{r=1}^{t}r\beta^{\chi(n-\phi)}C_1 \qquad \text{by (15) and (13)}\\ &\le -\frac{t^2}{2}\beta^{\chi(n-\phi)}C_1.\end{aligned} \tag{21}$$

Therefore, by (19), (20) and (21), almost surely, the following hold

$$\lim_{t\to\infty}\frac{1}{t^2}\psi_t^i(\theta,\theta^*) \le -\frac{1}{2}\beta^{\chi(n-\phi)}C_1.$$

Therefore, we have $\psi_t^i(\theta,\theta^*) \xrightarrow{\text{a.s.}} -\infty$ and $\mu_t^i(\theta) \xrightarrow{\text{a.s.}} 0$ for $i \in \mathcal{N}$ and $\theta \neq \theta^*$, proving Theorem 5. □

We now present the proof of our key lemma – Lemma 3.

Proof (Proof of Lemma 3). By (9), we have

$$\left|\mathcal{L}_r^i(\theta,\theta^*)\right| = \left|\log\frac{\ell_i(s_t^i|\theta)}{\ell_i(s_t^i|\theta^*)}\right| \le \max_{i\in\mathcal{V}}\max_{\theta_1,\theta_2\in\Theta;\theta_1\neq\theta_2}\max_{w_i\in\mathcal{S}_i}\left|\log\frac{\ell_i(w_i|\theta_1)}{\ell_i(w_i|\theta_2)}\right|.$$

Note that $\max_{i\in\mathcal{V}}\max_{\theta_1,\theta_2\in\Theta;\theta_1\neq\theta_2}\max_{w_i\in\mathcal{S}_i}\left|\log\frac{\ell_i(w_i|\theta_1)}{\ell_i(w_i|\theta_2)}\right|$ is symmetric in θ_1 and θ_2. Thus,

$$\begin{aligned}\left|\mathcal{L}_r^i(\theta,\theta^*)\right| &\le \max_{i\in\mathcal{V}}\max_{\theta_1,\theta_2\in\Theta;\theta_1\neq\theta_2}\max_{w_i\in\mathcal{S}_i}\left|\log\frac{\ell_i(w_i|\theta_1)}{\ell_i(w_i|\theta_2)}\right| = \max_{i\in\mathcal{V}}\max_{\theta_1,\theta_2\in\Theta;\theta_1\neq\theta_2}\max_{w_i\in\mathcal{S}_i}\log\frac{\ell_i(w_i|\theta_1)}{\ell_i(w_i|\theta_2)}\\ &= \max_{i\in\mathcal{V}}\max_{\theta_1,\theta_2\in\Theta;\theta_1\neq\theta_2}\max_{w_i\in\mathcal{S}_i}-\log\frac{\ell_i(w_i|\theta_2)}{\ell_i(w_i|\theta_1)}\\ &= -\min_{i\in\mathcal{V}}\min_{\theta_1,\theta_2\in\Theta;\theta_1\neq\theta_2}\min_{w_i\in\mathcal{S}_i}\log\frac{\ell_i(w_i|\theta_2)}{\ell_i(w_i|\theta_1)} = -(-C_0) = C_0 < \infty. \end{aligned} \tag{22}$$

Thus, adding and subtracting $\frac{1}{t^2}\sum_{r=1}^{t}\sum_{j=1}^{n-\phi}\pi_j(r+1)\sum_{k=1}^{r}\mathcal{L}_k^j(\theta,\theta^*)$ from the first term on the right hand side of (19), we can get

$$\begin{aligned}&\frac{1}{t^2}\sum_{r=1}^{t}\left(\sum_{j=1}^{n-\phi}\mathbf{\Phi}_{ij}(t,r+1)\sum_{k=1}^{r}\mathcal{L}_k^j(\theta,\theta^*) - \pi_j(r+1)r\sum_{j=1}^{n-\phi}H_j(\theta,\theta^*)\right)\\ &= \frac{1}{t^2}\sum_{r=1}^{t}\sum_{j=1}^{n-\phi}\left(\mathbf{\Phi}_{ij}(t,r+1)-\pi_j(r+1)\right)\sum_{k=1}^{r}\mathcal{L}_k^j(\theta,\theta^*)\\ &\quad+\frac{1}{t^2}\sum_{r=1}^{t}\sum_{j=1}^{n-\phi}\pi_j(r+1)\left(\sum_{k=1}^{r}\mathcal{L}_k^j(\theta,\theta^*) - rH_j(\theta,\theta^*)\right).\end{aligned} \tag{23}$$

For the first term of the right hand side of (23), we have

$$\begin{aligned}&\frac{1}{t^2}\left|\sum_{r=1}^{t}\sum_{j=1}^{n-\phi}\left(\mathbf{\Phi}_{ij}(t,r+1)-\pi_j(r+1)\right)\sum_{k=1}^{r}\mathcal{L}_k^j(\theta,\theta^*)\right|\\ &\le \frac{1}{t^2}\sum_{r=1}^{t}\sum_{j=1}^{n-\phi}\left|\mathbf{\Phi}_{ij}(t,r+1)-\pi_j(r+1)\right|\sum_{k=1}^{r}\left|\mathcal{L}_k^j(\theta,\theta^*)\right| \qquad (24)\\ &\le \frac{1}{t^2}\sum_{r=1}^{t}\sum_{j=1}^{n-\phi}\left|\mathbf{\Phi}_{ij}(t,r+1)-\pi_j(r+1)\right| rC_0 \quad \text{by (22)}\\ &\le \frac{1}{t^2}\sum_{r=1}^{t}\sum_{j=1}^{n-\phi}(1-\beta^\nu)^{\lceil\frac{t-r}{\nu}\rceil}rC_0 \quad \text{by Theorem 3}\\ &\le \frac{1}{t^2}\left(t(n-\phi)C_0\right)\sum_{r=1}^{t}(1-\beta^\nu)^{\lceil\frac{t-r}{\nu}\rceil}\\ &\le \frac{(n-\phi)C_0}{(1-\beta^\nu)(1-(1-\beta^\nu)^{\frac{1}{\nu}})t}. \qquad (25)\end{aligned}$$

Thus, for every sample path, we have

$$\frac{1}{t^2}\sum_{r=1}^{t}\sum_{j=1}^{n-\phi}\left(\mathbf{\Phi}_{ij}(t,r+1)-\pi_j(r+1)\right)\sum_{k=1}^{r}\mathcal{L}_k^j(\theta,\theta^*)\to 0.$$

For the second term of the right hand side of (23), we will show that

$$\frac{1}{t^2}\sum_{r=1}^{t}\sum_{j=1}^{n-\phi}\pi_j(r+1)\left(\sum_{k=1}^{r}\mathcal{L}_k^j(\theta,\theta^*)-rH_j(\theta,\theta^*)\right)\xrightarrow{\text{a.s.}}0,$$

i.e., almost surely for any $\epsilon>0$ there exists sufficiently large $t(\epsilon)$ such that $\forall\, t\geq t(\epsilon)$,

$$\frac{1}{t^2}\left|\sum_{r=1}^{t}\sum_{j=1}^{n-\phi}\pi_j(r+1)\left(\sum_{k=1}^{r}\mathcal{L}_k^j(\theta,\theta^*)-rH_j(\theta,\theta^*)\right)\right| \;\leq\; \epsilon. \tag{26}$$

We prove this by dividing r into two ranges $r\in\{1,\cdots,\sqrt{t}\}$ and $r\in\{\sqrt{t}+1,\cdots,t\}$, i.e.,

$$\begin{aligned}
&\frac{1}{t^2}\sum_{r=1}^{t}\sum_{j=1}^{n-\phi}\pi_j(r+1)\left(\sum_{k=1}^{r}\mathcal{L}_k^j(\theta,\theta^*)-rH_j(\theta,\theta^*)\right)\\
&=\frac{1}{t^2}\sum_{r=1}^{\sqrt{t}}\sum_{j=1}^{n-\phi}\pi_j(r+1)\left(\sum_{k=1}^{r}\mathcal{L}_k^j(\theta,\theta^*)-rH_j(\theta,\theta^*)\right)\\
&\quad+\frac{1}{t^2}\sum_{r=\sqrt{t}+1}^{t}\sum_{j=1}^{n-\phi}\pi_j(r+1)\left(\sum_{k=1}^{r}\mathcal{L}_k^j(\theta,\theta^*)-rH_j(\theta,\theta^*)\right).
\end{aligned}\tag{27}$$

For the first term of the right hand side of (27), we have

$$\begin{aligned}
&\frac{1}{t^2}\left|\sum_{r=1}^{\sqrt{t}}\sum_{j=1}^{n-\phi}\pi_j(r+1)\left(\sum_{k=1}^{r}\mathcal{L}_k^j(\theta,\theta^*)-rH_j(\theta,\theta^*)\right)\right|\\
&\leq\frac{1}{t^2}\sum_{r=1}^{\sqrt{t}}\sum_{j=1}^{n-\phi}\pi_j(r+1)\,(2rC_0)\quad\text{by (13) and (22)}\\
&=\frac{1}{t^2}\,(2C_0)\sum_{r=1}^{\sqrt{t}}r\\
&\leq C_0\left(\frac{1}{t}+\frac{1}{t^{\frac{3}{2}}}\right).
\end{aligned}$$

Thus, there exists $t_1(\epsilon)$ such that for all $t\geq t_1(\epsilon)$, it holds that

$$\frac{1}{t^2}\left|\sum_{r=1}^{\sqrt{t}}\sum_{j=1}^{n-\phi}\pi_j(r+1)\left(\sum_{k=1}^{r}\mathcal{L}_k^j(\theta,\theta^*)-rH_j(\theta,\theta^*)\right)\right|\leq\frac{\epsilon}{2}.$$

For the second term of the right hand side of (27), we have

$$\begin{aligned}
&\frac{1}{t^2}\sum_{r=\sqrt{t}+1}^{t}\sum_{j=1}^{n-\phi}\pi_j(r+1)\left(\sum_{k=1}^{r}\mathcal{L}_k^j(\theta,\theta^*)-rH_j(\theta,\theta^*)\right)\\
&=\frac{1}{t}\sum_{r=\sqrt{t}+1}^{t}\sum_{j=1}^{n-\phi}\pi_j(r+1)\frac{r}{t}\left(\frac{1}{r}\sum_{k=1}^{r}\mathcal{L}_k^j(\theta,\theta^*)-H_j(\theta,\theta^*)\right)
\end{aligned}$$

Since $\mathcal{L}_k^j(\theta,\theta^*)$'s are i.i.d., from Strong LLN, we know that $\frac{1}{r}\sum_{k=1}^{r}\mathcal{L}_k^j(\theta,\theta^*) - H_j(\theta,\theta^*) \xrightarrow{\text{a.s.}} 0$. That is, with probability 1, the sample path converges. Now, focus on each convergent sample path. For sufficiently large $r(\epsilon)$, it holds that for any $r \geq r(\epsilon)$,

$$\left|\frac{1}{r}\sum_{k=1}^{r}\mathcal{L}_k^j(\theta,\theta^*) - H_j(\theta,\theta^*)\right| \leq \frac{\epsilon}{2}.$$

Recall that $r \geq \sqrt{t}$. Thus, we know that there exists sufficiently large $t_2(\epsilon)$ such that $\forall\, t \geq t_2(\epsilon)$, $r \geq \sqrt{t}$ is large enough and

$$\left|\frac{1}{r}\sum_{k=1}^{r}\mathcal{L}_k^j(\theta,\theta^*) - H_j(\theta,\theta^*)\right| \leq \frac{\epsilon}{2}.$$

Then, we have $\forall\, t \geq t_2(\epsilon)$,

$$\begin{aligned}
&\frac{1}{t^2}\left|\sum_{r=\sqrt{t}+1}^{t}\sum_{j=1}^{n-\phi}\pi_j(r+1)\left(\sum_{k=1}^{r}\mathcal{L}_k^j(\theta,\theta^*) - rH_j(\theta,\theta^*)\right)\right| \\
&= \frac{1}{t}\sum_{r=\sqrt{t}+1}^{t}\sum_{j=1}^{n-\phi}\pi_j(r+1)\frac{r}{t}\left|\frac{1}{r}\sum_{k=1}^{r}\mathcal{L}_k^j(\theta,\theta^*) - H_j(\theta,\theta^*)\right| \\
&\leq \frac{1}{t}\sum_{r=\sqrt{t}+1}^{t}\sum_{j=1}^{n-\phi}\pi_j(r+1)\frac{r}{t}\frac{\epsilon}{2} \\
&= \frac{1}{t}\sum_{r=\sqrt{t}+1}^{t}\frac{r}{t}\frac{\epsilon}{2} = \frac{\epsilon}{2}\frac{1}{t^2}\sum_{r=\sqrt{t}+1}^{t} r \\
&= \frac{\epsilon}{4}\frac{1}{t^2}\left(t^2-\sqrt{t}\right) \leq \frac{\epsilon}{2}.
\end{aligned}$$

Therefore, for any $\epsilon > 0$, there exists $\max\{t_1(\epsilon), t_2(\epsilon)\}$, such that for any $t \geq \max\{t_1(\epsilon), t_2(\epsilon)\}$,

$$\frac{1}{t^2}\left|\sum_{r=1}^{t}\sum_{j=1}^{n-\phi}\pi_j(r+1)\left(\sum_{k=1}^{r}\mathcal{L}_k^j(\theta,\theta^*) - rH_j(\theta,\theta^*)\right)\right| \leq \epsilon,$$

for every convergent sample path. In addition, we know a sample path is convergent with probability 1. Thus (26) holds almost surely.

Therefore, Lemma 3 is proved.

□

5 BFL in the absence of Byzantine Agents, i.e., $f = 0$

In this section, we present BFL for the special case in the absence of Byzantine agents, i.e., $f = 0$, named Failure-free BFL. Since $f = 0$, all the agents in the network are cooperative, and no trimming is needed. Indeed, the BFL for $f = 0$ is a simple modification of the algorithm proposed in [14].

For each time $t \geq 1$, we define a matrix that follows the structure of $G(\mathcal{V},\mathcal{E})$ as follows:

$$\mathbf{A}_{ij} \triangleq \begin{cases} \frac{1}{|\mathcal{I}_i|+1}, & j \in \mathcal{I}_i \cup \{i\} \\ 0, & \text{otherwise.} \end{cases} \tag{28}$$

Algorithm 5: Failure-free BFL

1 Transmit current belief vector μ_{t-1}^i on all outgoing edges;
2 Wait until a private signal s_t^i is observed and belief vectors are received from all incoming neighbors $\mathcal{I}_i$;
3 **for** $\theta \in \Theta$ **do**
4 $\quad \mu_t^i(\theta) \leftarrow \dfrac{\ell_i(s_{1,t}^i|\theta)\prod_{j\in\mathcal{I}_i\cup\{i\}}\mu_{t-1}^j(\theta)^{\frac{1}{|\mathcal{I}_i|+1}}}{\sum_{p=1}^m \ell_i(s_{1,t}^i|\theta)\prod_{j\in\mathcal{I}_i\cup\{i\}}\mu_{t-1}^j(\theta)^{\frac{1}{|\mathcal{I}_i|+1}}}$.
5 **end**

Thus, the dynamic of $\psi_t^i(\theta,\theta^*)$ (defined in (9)) under Algorithm 5 can be written as

$$\begin{aligned}
\psi_t^i(\theta,\theta^*) &= \log\frac{\mu_t^i(\theta)}{\mu_t^i(\theta^*)}\\
&= \log\frac{\ell_i(s_{1,t}^i|\theta)\prod_{j\in\mathcal{I}_i\cup\{i\}}\mu_{t-1}^j(\theta)^{\frac{1}{|\mathcal{I}_i|+1}}}{\ell_i(s_{1,t}^i|\theta^*)\prod_{j\in\mathcal{I}_i\cup\{i\}}\mu_{t-1}^j(\theta^*)^{\frac{1}{|\mathcal{I}_i|+1}}}\\
&= \log\prod_{j\in\mathcal{I}_i\cup\{i\}}\left[\frac{\mu_{t-1}^j(\theta)}{\mu_{t-1}^j(\theta^*)}\right]^{\frac{1}{|\mathcal{I}_i|+1}} + \log\frac{\ell_i(s_{1,t}^i\mid\theta)}{\ell_i(s_{1,t}^i\mid\theta^*)}\\
&= \log\prod_{j\in\mathcal{I}_i\cup\{i\}}\left[\frac{\mu_{t-1}^j(\theta)}{\mu_{t-1}^j(\theta^*)}\right]^{\frac{1}{|\mathcal{I}_i|+1}} + \sum_{r=1}^t\log\frac{\ell_i(s_r^i\mid\theta)}{\ell_i(s_r^i\mid\theta^*)}\\
&= \sum_{j=1}^n \mathbf{A}_{ij}\psi_{t-1}^i(\theta,\theta^*) + \sum_{r=1}^t\mathcal{L}_r^i(\theta,\theta^*) \quad \text{by (9) and (28)}
\end{aligned}$$

Recall that $\boldsymbol{\psi}_t(\theta,\theta^*)\in\mathbb{R}^{n-\phi}$ is the vector that stacks $\psi_{t-1}^i(\theta,\theta^*)$ with the i–th entry being $\psi_{t-1}^i(\theta,\theta^*)$ for all $i\in\mathcal{N}$. Since $f=0$, i.e., the network is free of failures, it holds that

$$0\le\phi=|\mathcal{F}|\le f=0.$$

Thus, $\boldsymbol{\psi}_t(\theta,\theta^*)\in\mathbb{R}^n$. Similar to (12), the evolution of $\boldsymbol{\psi}_t(\theta,\theta^*)$ can be compactly written as follows.

$$\begin{aligned}
\boldsymbol{\psi}_t(\theta,\theta^*) &= \mathbf{A}^t\boldsymbol{\psi}_0(\theta,\theta^*) + \sum_{r=1}^t\mathbf{A}^{t-r}\sum_{k=1}^r\mathcal{L}_k(\theta,\theta^*)\\
&= \sum_{r=1}^t\mathbf{A}^{t-r}\sum_{k=1}^r\mathcal{L}_k(\theta,\theta^*). \qquad (29)
\end{aligned}$$

The last equality holds from the fact that $\boldsymbol{\psi}_0(\theta,\theta^*)=\mathbf{0}$.

As mentioned before, the non-Bayesian learning rules [14,15,11,17] are consensus-based learning algorithms, wherein agents are required to reach a common decision asymptotically.

Assumption 3 *The underlying communication network $G(\mathcal{V},\mathcal{E})$ is strongly connected.*

It is easy to see that $G(\mathcal{V},\mathcal{E})$ itself is the only reduced graph of $G(\mathcal{V},\mathcal{E})$, and that Assumption 3 is the special case of Assumption 1 when $f=0$. Thus,

$$\chi_m=1, \qquad \text{and} \qquad \nu_m=\chi_m(n-\phi)=n.$$

Note that both χ_m and ν_m are independent of m when $f = 0$. Henceforth in this section, we drop the subscripts of χ_m and ν_m for ease of notation.

Similar to (3), for any $r \geq 1$, we get

$$\lim_{t \geq r,\ t \to \infty} \mathbf{A}^{t-r} = \mathbf{1}\boldsymbol{\pi}.$$

Since $\mathbf{A}$ is time-invariant, the product limit $\lim_{t \geq r,\ t \to \infty} \mathbf{A}^{t-r}$ is also independent of r.

It is easy to see that

$$\mathbf{A} \geq \frac{1}{n}\mathbf{H},$$

where $\mathbf{H}$ is the adjacency matrix of the communication graph $G(\mathcal{V}, \mathcal{E})$, and that

$$\pi_j \geq \frac{1}{n^n}, \quad \forall\, j = 1, \cdots, n. \tag{30}$$

The following corollary is a direct consequence of Theorem 3, and its proof is omitted.

Corollary 1. *For all $t \geq r \geq 1$, it holds that $\left|[\mathbf{A}^{t-r}]_{ij} - \pi_j\right| \leq (1 - \frac{1}{n^n})^{\lceil \frac{t-r}{n} \rceil}$, where $[\mathbf{A}^{t-r}]_{ij}$ is the i, j–th entry of matrix $\mathbf{A}^{t-r}$.*

In addition, when $f = 0$, Assumption 2 becomes

Assumption 4 *Suppose that Assumption 3 holds. For any $\theta \neq \theta^*$, the following holds*

$$\sum_{j=1}^{m} D\left(\ell_j(\cdot|\theta^*) \parallel \ell_j(\cdot|\theta)\right) \ \neq \ 0. \tag{31}$$

As an immediate consequence of Theorem 5, we have the following corollary.

Corollary 2. *When Assumption 4 holds, each agent i will concentrate its belief on the true hypothesis θ^* almost surely, i.e., $\mu_t^i(\theta) \xrightarrow{\text{a.s.}} 0$ for all $\theta \neq \theta^*$.*

Since Corollary 2 is the special case of Theorem 5 for $f = 0$, the proof of Corollary 2 is omitted.

5.1 Finite-Time Analysis of Failure-Free BFL

In this subsection, we present the convergence rate that is achievable in finite time with high probability. Our proof is similar to the proof presented in [14,17].

Lemma 4. *Let $\lambda \triangleq \left(1 - (\frac{1}{n})^n\right)^{\frac{1}{n}}$, and let $\theta \neq \theta^*$, and consider $\psi_t^i(\theta, \theta^*)$ as defined in (9). Then, for each agent i we have*

$$\mathbb{E}\left[\psi_t^i(\theta, \theta^*)\right] \leq \frac{nC_0}{(1 - \frac{1}{n^n})(1 - \lambda)} t - \frac{C_1}{2n^n} t^2.$$

Proof. By (29), we have $\psi_t^i(\theta,\theta^*) = \sum_{r=1}^t \sum_{j=1}^n [\mathbf{A}^{t-r}]_{ij} \sum_{k=1}^r \mathcal{L}_k^j(\theta,\theta^*)$. Taking expectation of $\psi_t^i(\theta,\theta^*)$ with respect to $\ell^i(\cdot \mid \theta^*)$, we get

$$\begin{aligned}
\mathbb{E}^*\left[\boldsymbol{\psi}_t^i(\theta,\theta^*)\right] &= \mathbb{E}^*\left[\sum_{r=1}^t \sum_{j=1}^n [\mathbf{A}^{t-r}]_{ij} \sum_{k=1}^r \mathcal{L}_k^j(\theta,\theta^*)\right] \\
&= \sum_{r=1}^t \sum_{j=1}^n [\mathbf{A}^{t-r}]_{ij} \sum_{k=1}^r \mathbb{E}^*\left[\mathcal{L}_k^j(\theta,\theta^*)\right] \\
&= \sum_{r=1}^t \sum_{j=1}^n [\mathbf{A}^{t-r}]_{ij} r H_j(\theta,\theta^*) \quad \text{by} \quad (13) \\
&= \sum_{r=1}^t \sum_{j=1}^n \left([\mathbf{A}^{t-r}]_{ij} - \pi_j\right) r H_j(\theta,\theta^*) + \sum_{r=1}^t \sum_{j=1}^n \pi_j r H_j(\theta,\theta^*). \qquad (32)
\end{aligned}$$

For the first term in the right hand side of (32), we have

$$\begin{aligned}
\sum_{r=1}^t \sum_{j=1}^n \left([\mathbf{A}^{t-r}]_{ij} - \pi_j\right) r H_j(\theta,\theta^*) &\le \sum_{r=1}^t \sum_{j=1}^n \left|[\mathbf{A}^{t-r}]_{ij} - \pi_j\right| r \left|H_j(\theta,\theta^*)\right| \\
&\le \sum_{r=1}^t \sum_{j=1}^n \left[1 - \frac{1}{n^n}\right]^{\lceil \frac{t-r}{n} \rceil} r C_0 \quad \text{by Corollary 1, and (14)} \\
&= n C_0 \sum_{r=1}^t \left[1 - \frac{1}{n^n}\right]^{\lceil \frac{t-r}{n} \rceil} r \\
&\le \frac{n C_0}{(1 - \frac{1}{n^n})(1-\lambda)} t. \qquad (33)
\end{aligned}$$

Since $G(\mathcal{V},\mathcal{E})$ is the only source component, C_1 (defined in (15)) becomes

$$C_1 = \min_{\theta,\theta^* \in \Theta; \theta \neq \theta^*} \sum_{i=1}^n D(\ell_i(\cdot|\theta^*) \parallel \ell_i(\cdot|\theta)).$$

Thus, for the second term in the right hand side of (32), we get

$$\begin{aligned}
\sum_{r=1}^t \sum_{j=1}^n \pi_j r H_j(\theta,\theta^*) &\le \sum_{r=1}^t \sum_{j=1}^n \frac{1}{n^n} r H_j(\theta,\theta^*) \quad \text{by (30) and (13)} \\
&= \frac{1}{n^n} \sum_{r=1}^t r \sum_{j=1}^n H_j(\theta,\theta^*) \\
&\le -\frac{1}{n^n} \sum_{r=1}^t r C_1 \\
&\le -\frac{C_1}{2n^n} t^2. \qquad (34)
\end{aligned}$$

By (33) and (34), (32) becomes

$$\begin{aligned}\mathbb{E}^*\left[\boldsymbol{\psi}_t^i(\theta,\theta^*)\right] &= \sum_{r=1}^{t}\sum_{j=1}^{n}\left([\mathbf{A}^{t-r}]_{ij}-\pi_j\right) rH_j(\theta,\theta^*) + \sum_{r=1}^{t}\sum_{j=1}^{n}\pi_j rH_j(\theta,\theta^*) \\ &\leq \frac{nC_0}{(1-\frac{1}{n^n})(1-\lambda)}t - \frac{C_1}{2n^n}t^2, \end{aligned} \tag{35}$$

proving the lemma.

□

Similar to [14,17], we also use McDiarmid's Inequality.

Theorem 6 (McDiarmid's Inequality). *Let $X_1,\cdots,X_t$ be independent random variables and consider the mapping $H:\mathcal{X}^t\to\mathbb{R}$. If for $r=1,\cdots,t$, and every sample $x_1,\cdots,x_t$, $x_r'\in\mathcal{X}$, the function H satisfies*

$$\left|H(x_1,\cdots,x_r,\cdots,x_t)-H(x_1,\cdots,x_r',\cdots,x_t)\right|\leq c_r,$$

then for all $\epsilon>0$,

$$\mathbb{P}\left[|H(x_1,\cdots,x_t)-\mathbb{E}[H(x_1,\cdots,x_t)]|\geq\epsilon\right]\leq\exp\left\{\frac{-2\epsilon^2}{\sum_{r=1}^{t}c_r^2}\right\}.$$

Theorem 7. *Under Assumption 4, for any $\rho\in(0,1)$, there exists an integer $T(\rho)$ such that with probability $1-\rho$, for all $t\geq T(\rho)$ and for all $\theta\neq\theta^*$, we have*

$$\mu_t^i(\theta)\leq\exp\left(\frac{nC_0}{(1-\frac{1}{n^n})(1-\lambda)}t-\frac{C_1}{4n^n}t^2\right)$$

where C_0 and C_1 are defined in (14) and (15) respectively, and $T(\rho)=\frac{64C_0^2n^{2n}}{3C_1^2}\log\frac{1}{\rho}$.

Proof. Since $\mu_t^i(\theta^*)\in(0,1]$, we have

$$\mu_t^i(\theta)\leq\frac{\mu_t^i(\theta)}{\mu_t^i(\theta^*)}=\exp\left(\psi_t^i(\theta,\theta^*)\right).$$

Thus, we have

$$\begin{aligned}\mathbb{P}\left(\mu_t^i(\theta)\geq\exp\left(\frac{nC_0}{(1-\frac{1}{n^n})(1-\lambda)}t-\frac{C_1}{4n^n}t^2\right)\right) &\leq \mathbb{P}\left(\psi_t^i(\theta,\theta^*)\geq\frac{nC_0}{(1-\frac{1}{n^n})(1-\lambda)}t-\frac{C_1}{4n^n}t^2\right) \\ &\leq \mathbb{P}\left(\psi_t^i(\theta,\theta^*)-\mathbb{E}^*\left[\psi_t^i(\theta,\theta^*)\right]\geq\frac{C_1}{4n^n}t^2\right).\end{aligned}$$

Note that $\psi_t^i(\theta,\theta^*)$ is a function of the random vector $\mathbf{s}_1,\cdots,\mathbf{s}_t$. For a given sample path $\mathbf{s}_1,\cdots,\mathbf{s}_t$, and for all $p\in\{1,\cdots,t\}$, we have

$$\begin{aligned}
&\max_{\mathbf{s}_p\in\mathcal{S}_1\times\cdots\times\mathcal{S}_t}\psi_t^i(\theta,\theta^*)-\min_{\mathbf{s}_p\in\mathcal{S}_1\times\cdots\times\mathcal{S}_t}\psi_t^i(\theta,\theta^*)\\
&=\max_{\mathbf{s}_p\in\mathcal{S}_1\times\cdots\times\mathcal{S}_t}\sum_{r=1}^{t}\sum_{j=1}^{n}[\mathbf{A}^{t-r}]_{ij}\sum_{k=1}^{r}\mathcal{L}_k(\theta,\theta^*)-\min_{\mathbf{s}_p\in\mathcal{S}_1\times\cdots\times\mathcal{S}_t}\sum_{r=1}^{t}\sum_{j=1}^{n}[\mathbf{A}^{t-r}]_{ij}\sum_{k=1}^{r}\mathcal{L}_k(\theta,\theta^*)\\
&=\max_{\mathbf{s}_p\in\mathcal{S}_1\times\cdots\times\mathcal{S}_t}\sum_{r=p}^{t}\sum_{j=1}^{n}[\mathbf{A}^{t-r}]_{ij}\sum_{k=1}^{r}\mathcal{L}_k(\theta,\theta^*)-\min_{\mathbf{s}_p\in\mathcal{S}}\sum_{r=p}^{t}\sum_{j=1}^{n}[\mathbf{A}^{t-r}]_{ij}\sum_{k=1}^{r}\mathcal{L}_k(\theta,\theta^*)\\
&=\max_{\mathbf{s}_p\in\mathcal{S}_1\times\cdots\times\mathcal{S}_t}\sum_{r=p}^{t}\sum_{j=1}^{n}[\mathbf{A}^{t-r}]_{ij}\mathcal{L}_p(\theta,\theta^*)-\min_{\mathbf{s}_p\in\mathcal{S}}\sum_{r=p}^{t}\sum_{j=1}^{n}[\mathbf{A}^{t-r}]_{ij}\mathcal{L}_p(\theta,\theta^*)\\
&\le\sum_{r=p}^{t}\sum_{j=1}^{n}[\mathbf{A}^{t-r}]_{ij}C_0+\sum_{r=p}^{t}\sum_{j=1}^{n}[\mathbf{A}^{t-r}]_{ij}C_0\\
&=2C_0(t-p+1)\triangleq c_p.
\end{aligned}$$

By McDiarmid's inequality (Theorem 6), we obtain that

$$\begin{aligned}
\mathbb{P}\left(\psi_t^*(\theta,\theta^*)-\mathbb{E}^*\left[\psi_t^*(\theta,\theta^*)\right]\ge\frac{C_1}{4n^n}t^2\right)&\le\exp\left(-\frac{2\frac{C_1^2}{16n^{2n}}t^4}{\sum_{p=1}^{t}(2C_0(t-p+1))^2}\right)\\
&\le\exp\left(-\frac{3C_1^2}{64C_0^2n^{2n}}t\right),
\end{aligned}$$

where the last inequality follows from the fact that

$$t(t+1)(2t+1)\le 4t^3\quad\forall\, t\ge 2,$$

which can be shown by induction.

Therefore, for a given confidence level ρ, in order to have

$$\mathbb{P}\left(\mu_t^i(\theta)\ge\exp\left(\frac{nC_0}{(1-\frac{1}{n^n})(1-\lambda)}t-\frac{C_1}{4n^n}t^2\right)\right)\le\rho,$$

we require that

$$t\ge T(\rho)=\frac{64C_0^2n^{2n}}{3C_1^2}\log\frac{1}{\rho}.$$

□

Remark 3. The above finite-time analysis is not directly applicable for the general case when $f>0$, due to the fact that the local beliefs are dependent on all the observations collected so far *as well as* all the future observations.

Remark 4. Our analysis for the special when $f=0$ also works for time-varying networks [14]. In addition, with identical analysis, we are able to adapt the failure-free scheme to work in the more general setting where there is no underlying true state, and the goal is to have the agents collaboratively identify an optimal $\theta\in\Theta$ that best explains all the observations obtained over the whole network.

6 Modified BFL and Minimal Network Identifiability

To reduce the computation complexity per iteration in general, and to identify the minimal (tight) global identifiability of the network for any consensus-based learning rule of interest to learn the true state, we propose a modification of the above learning rule, which works under much weaker network topology and global identifiability condition.

We decompose the m-ary hypothesis testing problem into $m(m-1)$ (ordered) binary hypothesis testing problems. For each pair of hypotheses θ_1 and θ_2, each non-faulty agent updates the likelihood ratio of θ_1 over θ_2 as follows. Let $r_t^i(\theta_1, \theta_2)$ be the log likelihood ratio of θ_1 over θ_2 kept by agent i at the end of iteration t. Our modified learning rule applies consensus procedures to log likelihood ratio, i.e., $r_t^i(\theta_1, \theta_2)$, which is a scalar. For Algorithm 6, we only require scalar iterative Byzantine (approximate) consensus among the non-faulty agents to be achievable.

When scalar consensus is achievable, the following assumption on the identifiability of the network to detect θ^* is minimal, meaning that if this assumption is not satisfied, then no correct consensus-based non-Bayesian learning exists.

Assumption 5 *Suppose that every 1-dimensional reduced graph of $G(\mathcal{V}, \mathcal{E})$ contains only one source component. For any $\theta \neq \theta^*$, and for any 1-dimensional reduced graph $\mathcal{H}_1$ of $G(\mathcal{V}, \mathcal{E})$ with $\mathcal{S}_{\mathcal{H}_1}$ denoting the unique source component, the following holds*

$$\sum_{j \in \mathcal{S}_{\mathcal{H}_1}} D\left(\ell_j(\cdot|\theta^*) \parallel \ell_j(\cdot|\theta)\right) \neq 0. \tag{36}$$

Assumption 5 is minimal for the following reasons: (1) For any consensus-based learning rule to work, the communication network $G(\mathcal{V}, \mathcal{E})$ should support consensus with scalar inputs. That is, every 1-dimensional reduced graph of $G(\mathcal{V}, \mathcal{E})$ must contain only one source component. (2) Under some faulty behaviors of the Byzantine agents, one particular 1–dimensional reduced graph may govern the entire dynamics of $r^i(\theta_1, \theta_2)$. If (36) does not hold for that reduced graph, then the good agents may not able to distinguish θ_1 from θ_2.

Algorithm 6: Pairwise Learning

Initialization: **for** $\theta_1, \theta_2 \in \Theta$, *and* $\theta_1 \neq \theta_2$ **do**
  $r_0^i(\theta_1, \theta_2) \leftarrow 0$;
end
while $t \geq 1$ **do**
  for $\theta_1, \theta_2 \in \Theta$, *and* $\theta_1 \neq \theta_2$ **do**
    Transmit current belief vector $r_{t-1}^i(\theta_1, \theta_2)$ on all outgoing edges;
    Wait until a private signal s_t^i is observed and log likelihood ratios $\tilde{r}_{t-1}^j(\theta_1, \theta_2)$ are received from all incoming neighbors $\mathcal{I}_i$;
    Sort the received log likelihood ratios $\tilde{r}_{t-1}^j(\theta_1, \theta_2)$ in a non-decreasing order, and remove the smallest f values and the largest f values. % Denote the set of indices of incoming neighbors whose ratios have not been removed at iteration t by $\mathcal{I}_i^*[t]$.%
    $r_t^i(\theta_1, \theta_2) \leftarrow \frac{\sum_{j \in \mathcal{I}_i^*[t]} \tilde{r}_{t-1}^j(\theta_1, \theta_2) + r_{t-1}^i(\theta_1, \theta_2)}{|\mathcal{I}^*[t]|+1} + \log \frac{\ell_i(s_{1,t}^i|\theta_1)}{\ell_i(s_{1,t}^i|\theta_2)}$.
  end
end

For each iteration, the computation complexity per agent (non-faulty) can be calculated as follows. The cost-dominant procedure in each iteration is sorting the received log likelihood ratios, which takes $O(n \log n)$ operations. In total, we have $m(m-1)$ order pairs of hypotheses. Thus, the total computation per agent per iteration is $O(m^2 n \log n)$.

Theorem 8. *Suppose Assumption 5 holds. Under Algorithm 6, for any $\theta \neq \theta^*$, the following holds:*

$$r_t^i(\theta^*, \theta) \xrightarrow{\text{a.s.}} +\infty, \text{ and } \quad r_t^i(\theta, \theta^*) \xrightarrow{\text{a.s.}} -\infty.$$

Proof. By [20], we know that for each pair of hypotheses θ_1 and θ_2, there exists a row-stochastic matrix $\mathbf{M}^{1,2}[t] \in \mathbb{R}^{(n-\phi)\times(n-\phi)}$ such that

$$r_t^i(\theta_1, \theta_2) = \sum_{j=1}^{n-\phi} \mathbf{M}_{ij}^{1,2}[t] r_{t-1}^j(\theta_1, \theta_2) + \log \frac{\ell_i(s_{1,t}^i \mid \theta_1)}{\ell_i(s_{1,t}^i \mid \theta_2)}. \tag{37}$$

Note that matrix $\mathbf{M}^{1,2}$ depends on the choice of hypotheses θ_1 and θ_2.

For a given pair of hypotheses θ_1 and θ_2, let $\mathbf{r}_t(\theta_1, \theta_2) \in \mathbb{R}^{n-\phi}$ be the vector that stacks $r_t^i(\theta_1, \theta_2)$. The evolution of $\mathbf{r}(\theta_1, \theta_2)$ can be compactly written as

$$\begin{aligned} \mathbf{r}_t(\theta_1, \theta_2) &= \mathbf{M}^{1,2}[t] \mathbf{r}_{t-1}(\theta_1, \theta_2) + \sum_{r=1}^{t} \mathcal{L}_r(\theta_1, \theta_2) \\ &= \sum_{r=1}^{t} \mathbf{\Phi}^{1,2}(t, r+1) \sum_{k=1}^{r} \mathcal{L}_k(\theta_1, \theta_2), \end{aligned} \tag{38}$$

where $\mathbf{\Phi}^{1,2}(t, r+1) \triangleq \mathbf{M}^{1,2}[t]\mathbf{M}^{1,2}[t-1] \cdots \mathbf{M}^{1,2}[r+1]$ for $r \leq t$, $\mathbf{\Phi}^{1,2}(t,t) \triangleq \mathbf{M}^{1,2}[t]$ and $\mathbf{\Phi}^{1,2}(t, t+1) \triangleq \mathbf{I}$. We do the analysis for each pair of θ_1 and θ_2 separately.

The remaining proof is identical to the proof of Theorem 5, and is omitted.

□

Proposition 1. *Suppose there exists $\tilde{\theta} \in \Theta$ such that for any $\theta \neq \tilde{\theta}$, it holds that $r_t^i(\tilde{\theta}, \theta) \xrightarrow{\text{a.s.}} +\infty$, and $r_t^i(\theta, \tilde{\theta}) \xrightarrow{\text{a.s.}} -\infty$. Then $\tilde{\theta} = \theta^*$.*

Proof. We prove this proposition by contradiction. Suppose there exists $\tilde{\theta} \neq \theta^* \in \Theta$ such that for any $\theta \neq \tilde{\theta}$, it holds that $r_t^i(\tilde{\theta}, \theta) \xrightarrow{\text{a.s.}} +\infty$, and $r_t^i(\theta, \tilde{\theta}) \xrightarrow{\text{a.s.}} -\infty$. Then we know that $r_t^i(\tilde{\theta}, \theta^*) \xrightarrow{\text{a.s.}} +\infty$ and $r_t^i(\theta^*, \tilde{\theta}) \xrightarrow{\text{a.s.}} -\infty$, contradicting Theorem 8. Thus, Proposition 1 is true.

□

7 Conclusion

This paper addresses the problem of consensus-based non-Bayesian learning over multi-agent networks when an unknown subset of agents may be adversarial (Byzantine). We propose two learning rules, and characterize the tight network identifiability condition for any consensus-based learning rule of interest to exist. In our first update rule, each agent updates its local beliefs as (up to normalization) the product of (1) the likelihood of the *cumulative* private signals and (2) the weighted geometric average of the beliefs of its incoming neighbors and itself. Under reasonable assumptions on the underlying network structure and the global identifiability of the network, we show that all the non-faulty agents asymptotically agree on the true state almost surely. For the case when every agent is failure-free, we show that (with high probability) each agent's beliefs on the wrong

hypotheses decrease at rate $O(\exp(-Ct^2))$, where t is the number of iterations, and C is a constant. In general when agents may be adversarial, network identifiability condition specified for the above learning rule scales poorly in m. In addition, the computation complexity per agent per iteration of this learning rule is forbiddingly high. Thus, we propose a modification of our first learning rule, whose complexity per iteration per agent is $O(m^2 n \log n)$. We show that this modified learning rule works under a much weaker global identifiability condition that is independent of m.

We so far focussed on synchronous system and static network, our results may be generalizable to asynchronous as well as time varying network.

Throughout this paper, we assume that consensus among non-faulty agents needs to be achieved. Although this is necessary for the family of consensus-based algorithms (by definition), this is not the case for the non-faulty agents to collaboratively learn the true state in general. Indeed, there is a tradeoff between the capability of the network to reach consensus and the tight condition of the network detectability. For instance, if the network is disconnected, then information cannot be propagated across the connected components. Thus, the non-faulty agents in each connected component have to be able to learn the true state. We leave investigating the above tradeoff as future work.

References

1. J.-F. Chamberland and V. V. Veeravalli. Decentralized detection in sensor networks. *IEEE Transactions on Signal Processing*, 2003.
2. D. Gale and S. Kariv. Bayesian learning in social networks. *Games and Economic Behavior*, 2003.
3. A. Jadbabaie, P. Molavi, A. Sandroni, and A. Tahbaz-Salehi. Non-bayesian social learning. *Games and Economic Behavior*, 2012.
4. L. Su and N. H. Vaidya. Reaching approximate byzantine consensus with multi-hop communication. In *Proceedings of International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS)*, Aug. 2015.
5. J. Tsitsiklis. Decentralized detection by a large number of sensors. *Mathematics of Control, Signals and Systems*, 1988.
6. J. N. Tsitsiklis. Decentralized detection. *Advances in Statistical Signal Processing*, 1993.
7. P. K. Varshney. *Distributed Detection and Data Fusion*. Springer Science & Business Media, 2012.
8. E. Wong and B. Hajek. *Stochastic Processes in Engineering Systems*. Springer Science & Business Media, 2012.
9. J. Hajnal and M. Bartlett. Weak ergodicity in non-homogeneous markov chains. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 54, pages 233–246. Cambridge Univ Press, 1958.
10. A. Jadbabaie, P. Molavi, and A. Tahbaz-Salehi. Information heterogeneity and the speed of learning in social networks. *Columbia Business School Research Paper*, (13-28), 2013.
11. A. Lalitha, A. Sarwate, and T. Javidi. Social learning and distributed hypothesis testing. In *Information Theory (ISIT), 2014 IEEE International Symposium on*, pages 551–555, June 2014.
12. N. A. Lynch. *Distributed Algorithms*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1996.
13. P. Molavi and A. Jadbabaie. Foundations of non-bayesian social learning. *Columbia Business School Research Paper*, 2015.
14. A. Nedić, A. Olshevsky, and C. A. Uribe. Nonasymptotic convergence rates for cooperative learning over time-varying directed graphs. *arXiv preprint arXiv:1410.1977*, 2014.
15. K. R. Rad and A. Tahbaz-Salehi. Distributed parameter estimation in networks. In *IEEE Conference on Decision and Control (CDC)*, pages 5050–5055. IEEE, 2010.
16. S. Shahrampour and A. Jadbabaie. Exponentially fast parameter estimation in networks using distributed dual averaging. In *IEEE Conference on Decision and Control (CDC)*, pages 6196–6201. IEEE, 2013.
17. S. Shahrampour, A. Rakhlin, and A. Jadbabaie. Distributed detection: Finite-time analysis and impact of network topology. *arXiv preprint arXiv:1409.8606*, 2014.
18. S. Shahrampour, A. Rakhlin, and A. Jadbabaie. Finite-time analysis of the distributed detection problem. *arXiv preprint arXiv:1512.09311*, 2015.
19. J. Wolfowitz. Products of indecomposable, aperiodic, stochastic matrices. In *Proceedings of the American Mathematical Society*, pages 733–737. JSTOR, 1963.
20. N. H. Vaidya. Matrix Representation of Iterative Approximate Byzantine Consensus in Directed Graphs. arXiv 1203.1888, 2012.

21. D. Bajović, D. Jakovetić, J. M. Moura, J. Xavier, and B. Sinopoli. Large deviations performance of consensus+ innovations distributed detection with non-gaussian observations. *Signal Processing, IEEE Transactions on*, 60(11):5987–6002, 2012.
22. F. S. Cattivelli and A. H. Sayed. Distributed detection over adaptive networks using diffusion adaptation. *Signal Processing, IEEE Transactions on*, 59(5):1917–1932, 2011.
23. D. Jakovetić, J. M. Moura, and J. Xavier. Distributed detection over noisy networks: Large deviations analysis. *Signal Processing, IEEE Transactions on*, 60(8):4306–4320, 2012.
24. M. Pease ć R. Shostak ć L. Lamport. Reaching agreement in the presence of faults. J. ACM 27 (2), 228–234, Apr. 1980.
25. D. Dolev, N. A. Lynch, S. S. Pinter, E. W. Stark, and W. E. Weihl. Reaching approximate agreement in the presence of faults. *J. ACM*, 33(3):499–516, May 1986.
26. A. D. Fekete. Asymptotically optimal algorithms for approximate agreement. *Distributed Computing*, 4(1):9–29, 1990.
27. H. J. LeBlanc, H. Zhang, S. Sundaram, and X. Koutsoukos. Consensus of multi-agent networks in the presence of adversaries using only local information. In *Proceedings of the 1st International Conference on High Confidence Networked Systems*, HiCoNS '12, pages 1–10, New York, NY, USA, 2012. ACM.
28. N. H. Vaidya. Iterative byzantine vector consensus in incomplete graphs. In *Distributed Computing and Networking*, pages 14–28. Springer, 2014.
29. N. H. Vaidya and V. K. Garg. Byzantine vector consensus in complete graphs. In *Proceedings of the 2013 ACM symposium on Principles of distributed computing*, pages 65–73. ACM, 2013.
30. N. H. Vaidya, L. Tseng, and G. Liang. Iterative approximate byzantine consensus in arbitrary directed graphs. In *Proceedings of the 2012 ACM symposium on Principles of distributed computing*, pages 365–374. ACM, 2012.
31. H. Mendes and M. Herlihy. Multidimensional approximate agreement in byzantine asynchronous systems. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing*, STOC '13, pages 391–400, New York, NY, USA, 2013. ACM.
32. M.A. Perles and M. Sigron. A generalization of Tverberg's theorem. arXiv. 0710.4668, 2007.
33. A. Nedić, A. Olshevsky, and C. A. Uribe. Network Independent Rates in Distributed Learning arXiv preprint arXiv:1509.08574, 2015.
34. L. Su, and N. H. Vaidya. Asynchronous Distributed Hypothesis Testing in the Presence of Crash Failures University of Illinois at Urbana-Champaign, Tech. Rep, 2016. arXiv: 1606.03418.